

Envy Freedom and Prior-free Mechanism Design

Nikhil R. Devanur

Jason D. Hartline

Qiqi Yan

December 18, 2012

Abstract

We consider the provision of an abstract service to single-dimensional agents. Our model includes position auctions, single-minded combinatorial auctions, and constrained matching markets. When the agents' values are drawn from a distribution, the Bayesian optimal mechanism is given by Myerson (1981) as a virtual-surplus optimizer. We develop a framework for prior-free mechanism design and analysis. A good mechanism in our framework approximates the optimal mechanism for the distribution if there is a distribution; moreover, when there is no distribution this mechanism still performs well.

We define and characterize optimal envy-free outcomes in symmetric single-dimensional environments. Our characterization mirrors Myerson's theory. Furthermore, unlike in mechanism design where there is no point-wise optimal mechanism, there is always a point-wise optimal envy-free outcome.

Envy-free outcomes and incentive-compatible mechanisms are similar in structure and performance. We therefore use the optimal envy-free revenue as a benchmark for measuring the performance of a prior-free mechanism. A good mechanism is one that approximates the envy-free benchmark on any profile of agent values. We show that good mechanisms exist, and in particular, a natural generalization of the random sampling auction of Goldberg et al. (2001) is a constant approximation.

1 Introduction

The theories of optimal mechanism design for revenue maximization and social surplus maximization are fundamentally different. The revenue-optimal mechanism (Myerson, 1981) depends on the prior distribution from which the values of the agents are drawn whereas the surplus-optimal mechanism is prior-free (Vickrey, 1961; Clarke, 1971; Groves, 1973). In fact, the latter result is singular in this regard. Incentive constraints bind across possible misreports of each agent; therefore, an optimal mechanism must generally trade off performance on one input (i.e., profile of agent valuations) for another.

The goal of prior-free mechanism design and this work therein is to sacrifice optimality to obtain prior freedom. Of course, the hope is that not too much in the way of performance is lost. One way to make such a hope precise is to require that the prior-free mechanism perform within a constant factor of the Bayesian optimal mechanism when there is a distribution.

With this goal in mind, Hartline and Roughgarden (2008) give prior-free single-item auctions to maximize the agents' utility (surplus less payments). For any profile of valuations their mechanism approximates the performance of the best, in hindsight, (i.i.d.) Bayesian optimal mechanism. Approximation of this prior-free benchmark implies simultaneous approximation of the Bayesian optimal auction when the values of the agents are drawn from an i.i.d. distribution. Our main goal is

to extend this framework and approach to richer environments that include position auction (e.g., Varian, 2007; Edelman et al., 2007), single-minded combinatorial auctions (e.g., Lehmann et al., 2002), and constrained matching markets (e.g., Bikhchandani et al., 2008).

A first challenge with our goal is that the benchmark of Hartline and Roughgarden (2008) is analytically complex in general environments. Therefore, we consider an alternative approach motivated by the similarity between incentive and no-envy constraints. An outcome is envy free if no agent prefers the treatment of another to her own. I.e., no agent wishes to unilaterally swap outcomes with another. Incentive compatibility, on the other hand, requires that no agent wants to misreport her value. Importantly envy freedom, as a constraint, binds point wise on a valuation profile whereas incentive compatibility binds across different valuation profiles (for the possible misreports of the agents). While incentive-compatible mechanisms do not generally have point-wise optima, envy-free outcomes do.

A centerpiece of our work is a characterization of envy-free outcomes and their optima. This characterization is structurally equivalent to Myerson’s (1981) characterization of Bayesian optimal mechanisms applied to the empirical distribution given by the actual profile of agent values. The envy-free optimal outcome is a virtual-surplus maximizer. This connection between envy freedom and incentive compatibility is implicit in Jackson and Kremer (2007) where it is shown that they are equivalent in the limit, in contrast we show that they are structurally equivalent generally.

As a benchmark for prior-free mechanism design, the optimal envy-free performance is appealing as it inherits many nice properties of Bayesian optimal mechanisms. It is the maximum of a linear objective subject to feasibility constraints. Furthermore, in environments such as constrained matching markets, the envy-free optimal revenue point-wise dominates the revenue of the Bayesian-optimal incentive-compatible mechanism for any i.i.d. distribution. We view the envy-free optimal revenue as an unattainable prior-free “first-best” solution and, because of constraints of incentive compatibility, attempt to show that “second-best” approximation mechanisms are not too far away from it. While the Jackson and Kremer (2007) result implies such a result holds in the limit, we show that it continues to hold in general.

The incentive-compatible mechanisms that we design will satisfy the following prior-free performance guarantee: Their revenue on any profile of valuations will be within a constant multiplicative factor of the optimal envy-free revenue for the same valuation profile. The approximation factors we obtain range from 6.24 to roughly 2500. While these bounds are outside the realm of being practically significant in themselves, they do have important theoretical and practical implications. A mechanism that is a constant approximation in theory must exhibit at least some of the necessary characteristics needed for good performance in practice. As an example of the contrapositive, we show that mechanisms based on reserve prices do not give constant approximations in the environments that we consider. Furthermore, since our goal is theoretical understanding of good properties of mechanisms in environments where the optimal mechanism is complex, properties and analysis of constant approximations are provide more abstract understanding. Finally, the closer a mechanism is to optimal the more potential there is for the form of the mechanism to be overly dependent on modeling details that may not be accurate of practice. Hence, simple mechanisms that give constant approximations are often robust and detail free in the sense of Wilson (1987). Finally, in practice these mechanisms perform much better than their theoretical guarantees. For deeper motivation of approximation in mechanism design see the survey of Hartline (2011).

Our first approach for the design of prior-free mechanisms for general environments is via reduction. A digital-good environment (a.k.a., unlimited supply) is one where the mechanism has

no inter-agent feasibility constraint. A multi-unit environment (a.k.a., limited supply) is one where the mechanism has a constraint on the number of agents that can be simultaneously served, e.g., multiple units of a single item. We give a reduction from multi-unit environments to digital-good environments that loses at most a factor of two in approximation factor. We then give a lossless reduction from position auction environments and constrained matching environments to multi-unit environments. To obtain these reductions we give a structural characterization of these environments that shows that they are equivalent, even with respect to approximation. Given the 3.12-approximation for digital goods given by Ichiba and Iwama (2010), our reduction implies 6.24-approximations for multi-unit, constrained-matching, and position environments.

Our second approach is via a generalization of the random sampling auction of Goldberg et al. (2001) and Baliga and Vohra (2003). The auction takes the following form: The agents are randomly partitioned into a market and a sample. The sample is then used for market analysis and its empirical distribution is calculated. The optimal auction for the empirical distribution of the sample is then run on the market. This prior-free mechanism is one of the most fundamental, and we extend the analysis techniques derived for it in digital-good environments to multi-unit and (more generally) downward-closed environments. A downward-closed environment is given by a set system that specifies which agents can be simultaneously served. The only constraint placed on this set system is that subsets of feasible sets are themselves feasible. The approximation factors we obtain are 12.5 and 189 for multi-unit (and by the reduction above, constrained-matching and position auction environments) and downward-closed environments respectively.

Related Work. Our work fundamentally relies on the theory of optimal auctions as defined by Myerson (1981) and refined by Bulow and Roberts (1989). In particular, Myerson shows that Bayesian optimal mechanisms are virtual surplus optimizers and Bulow and Roberts show that the virtual value of an agent in this virtual surplus maximization can be viewed as the marginal revenue as given by an agent’s value distribution.

The random sampling auction for digital goods was first studied by Goldberg et al. (2001). The asymptotic performance of the mechanism was given by Segal (2003) and Baliga and Vohra (2003) and the convergence rate was studied by Balcan et al. (2008). In contrast, Goldberg et al. (2006a) consider the non-asymptotic behavior of the random sampling auction and show that its performance is a (large) constant factor from a prior-free benchmark that in retrospect coincides with ours. Alaei et al. (2009) give a nearly tight analysis that shows that the random sampling auction is a 4.68-approximation (the lower-bound is four).

There have been numerous attempts to design good prior-free mechanisms for digital goods outside the random sampling paradigm. Two notable approaches include an approximate reduction to the “decision problem”¹ by Goldberg et al. (2006a) and an approach based on statistical estimates that are non-manipulable with high probability by Goldberg and Hartline (2003). Hartline and McGrew (2005) extend the former approach and obtain an approximation factor of 3.25. Finally, Ichiba and Iwama (2010) show that a convex combination of these approaches gives an approximation factor of 3.12. The 6.24-approximation we obtain is the instantiation of our reduction with the 3.12 approximation of Ichiba and Iwama (2010).

The digital good auctions described above were analyzed in comparison to a natural single-priced benchmark. Hartline and Roughgarden (2008) suggest that approximation of a prior-free

¹Given a target profit, the decision problem is to construct a mechanism that obtains that target profit when it is attainable.

benchmark should imply approximation of the Bayesian optimal mechanism for any i.i.d. distribution. Benchmarks for which such an implication holds are well grounded in the classical theory of Bayesian optimal auctions. They propose the performance of the best, in hindsight, Bayesian optimal mechanism as a benchmark. For multi-unit environments, they characterize this benchmark as two-priced lotteries. In contrast, our benchmark, the envy-free optimal revenue, can be viewed as a relaxation of the Hartline-Roughgarden benchmark that is structurally simpler and analytically tractable in general downward-closed environments.

Subsequent to our study, Ha and Hartline (2012) generalized the statistical-estimation-based approach of Goldberg and Hartline (2003) to design a 32-approximation of the envy-free benchmark in downward-closed environments. A generalization of the reduction-to-the-decision-problem approach of Goldberg et al. (2006a) yields a 19-approximation (Ha and Hartline, 2011).

Overview. In Section 2 we describe in detail the our abstract environment for mechanism design which includes multi-unit, position, and combinatorial auctions. In Section 3 we characterize envy-free outcomes and their optima. In Section 4 we compare incentive-compatible and envy-free revenues. In Section 5 we describe our prior-free design and analysis framework and give Bayesian justification for using the envy-free optimal revenue as a prior-free performance benchmark. In Section 6 we describe a reduction-based approach wherein we show that, e.g., constrained matching mechanisms reduce to position auctions which reduce to multi-unit auctions which reduce (with a loss of a factor of two) to digital-good auctions. In Section 7 we describe a random sampling auction and analyze this auction in downward-closed environments.

2 Single-dimensional Environments

There are $n \geq 2$ agents. Each agent i has a valuation v_i for receiving an abstract service. The *valuation profile* is $\mathbf{v} = (v_1, \dots, v_n)$. We assume that the agents are indexed in order of decreasing values, i.e., $v_i \geq v_{i+1}$. An agent i who is served with probability x_i and charged price p_i obtains utility $u_i = v_i x_i - p_i$. *Individual rationality* requires that u_i be non-negative.

We are allowed to serve certain feasible sets of agents as given by a set system. The set system is *downward-closed* in the sense if a set of agents is feasible, so is any of its subsets. The empty set is always feasible. We allow randomization in two senses (1) the set system constraint may be randomized (i.e., by convex combination over set systems) and (2) the set of agents served may be random (by convex combination over feasible sets). Notably, randomization in (1) is given by the environment and randomization in (2) is by our choice of outcome. We define an *allocation* as a vector $\mathbf{x} = (x_1, \dots, x_n) \in [0, 1]^n$ where x_i is the probability that agent i is served. An allocation is feasible if it is the characteristic vector induced by the process above. The environments permitted include digital good auctions, multi-unit auctions, position auction environments, matroid environments, and single-minded combinatorial auctions.

We further assume that the feasibility constraint imposed by the environment is symmetric, i.e., the set of feasible allocations is closed under permutation. Digital good, multi-unit auction, and position auction environments are all symmetric. Given any asymmetric environment its corresponding *permutation environment* is obtained by randomly permuting the agents with respect to the feasibility constraint. Of special interest for us will be downward-closed permutation environments and matroid permutation environments. By definition these environments are symmetric.

To make the model described above precise, we give the following formal definitions.

Definition 2.1 (single-dimensional environment). *There are n agents denoted $N = \{1, \dots, n\}$. The sets of agents that can be simultaneously served are denoted by $\mathcal{X} \subset 2^N$. The mechanism may serve $S \subset N$ if $S \in \mathcal{X}$.*

Definition 2.2 (downward-closed environment). *A downward-closed environment is a single-dimensional environment where all subsets of feasible sets are feasible. I.e., $S \in \mathcal{X}$ and $T \subset S$ implies that $T \in \mathcal{X}$.*

Single-minded combinatorial auctions are an example of a downward-closed environment. In a single-minded combinatorial auction each agent i desires a specific bundle (i.e., subset) of a set of k available items. Agent i has value v_i for receiving her entire desired bundle (or a superset of it) and value zero otherwise. We say an agent is “served” if she receives her desired bundle. Notice that a set of agents S can be simultaneously served (i.e., $S \in \mathcal{X}$) if their desired bundles are disjoint. Of course, if S has disjoint desired bundles and $T \subset S$, then T has disjoint desired bundles, hence \mathcal{X} is downward closed.

Definition 2.3 (matroid environment). *A matroid environment is a downward-closed environment that satisfies an additional augmentation property: there is always an element in a larger cardinality feasible set that can be feasibly added to a smaller feasible set. I.e., given two feasible sets $S, T \subset \mathcal{X}$ with $|S| < |T|$, there exists an $i \in T \setminus S$ such that $S \cup \{i\} \in \mathcal{X}$.*

Two key consequences of the augmentation property are that (1) all maximal feasible sets have the same cardinality and that (2) the greedy algorithm optimizes surplus, i.e., the sum of the values of the agents served (see e.g., Oxley, 1992). The greedy algorithm sorts the agents by their values and then, in this order, greedily serves each agent if it is feasible to do so given the subset of agents previously served. Importantly the greedy algorithm is *ordinal* in that only the order of values matters in determining the surplus maximizing set and not the magnitude of the values. Matroid environments have a rich history in mechanism design, see e.g., Talwar (2003), Bikhchandani et al. (2008), and Hartline and Roughgarden (2009).

The example of a matroid environment that is most relevant to auction theory is that of *constrained matching markets*. In a constrained matching market each agent i desires one of a subset of k available items and her value for any item in this subset is v_i (her value for any other item is zero). The agent demand sets induce a bipartite graph between agents and items where an edge is between agent i and item j if j is one of i ’s desired items. A subset S of agents can be simultaneously served if there is a matching in the bipartite graph where all of S is matched. This matroid is known as the *transversal matroid*.

A special case of the transversal matroid is when the market is essentially partitioned into ℓ parts and within each part every item is acceptable to every agent (of course, there may be fewer items than agents in the part). This is known as the *partition matroid*. A special case of the partition matroid is the one where there is only one part, i.e., there are k identical items, and n agents who each desire one of these items. This matroid is known as the *k -uniform matroid* and it corresponds precisely to the standard k -unit auction environment.

Definition 2.4 (multi-unit environment). *There are n unit-demand agents and k units of an item. I.e., $\mathcal{X} = \{S : |S| \leq k\}$.*

An important special case of the multi-unit environment is one where the supply constraint k is never binding, i.e., when $k = n$. This special case is known as the *digital-good environment*.

Definition 2.5 (digital-good environment). *There are n unit-demand agents denoted $N = \{1, \dots, n\}$ and any subset of them can be served. I.e., $\mathcal{X} = 2^N$.*

A generalization of multi-unit environments that has recently been under intense scrutiny due to its application to auctions for selling advertisements on Internet search auction is the *position environment* (Varian, 2007; Edelman et al., 2007).

Definition 2.6 (position environment). *There are n agents and n positions. The positions have non-increasing weights $\mathbf{w} = (w_1, \dots, w_n)$. If an agent i is assigned position j she served with probability w_j and her value for this assignment is $v_i w_j$.*

It has been observed, e.g., by Dughmi et al. (2009), that there is a close connection between position environments and convex combinations of multi-unit environments. E.g., if we choose the supply k randomly so that k -units are available with probability $w_k - w_{k+1}$ then this distribution over k -unit environments gives the same service probabilities as the position auction. Of course, a k -unit auction is a position auction wherein the k highest weights are one, and the remaining weights are zero.

Digital-good, multi-unit, and position environments are agent-symmetric whereas matroid and downward-closed environments may not be. We can compose any environment with a permutation that renames the agents with respect to the set system. The convex combination of these environments with the permutation taken uniformly at random from the set of n -element permutations is referred to as a *permutation environment* and it is symmetric.

Definition 2.7. *Given any n agent single-dimensional environment specified by \mathcal{X} , the permutation environment for \mathcal{X} draws permutation π (mapping i to j by $\pi(i) = j$) uniformly at random from the set of all n -element permutations, and the feasible sets are given by $\mathcal{X}_\pi = \{\{\pi(i) : i \in S\} : S \in \mathcal{X}\}$.*

3 Optimal Envy-free Pricing

In this section we derive a theory of optimal envy-free outcomes in single-dimensional environments that mirrors that of Bayesian optimal (incentive compatible) mechanisms for i.i.d. prior distributions (cf., Myerson, 1981; Bulow and Roberts, 1989).

Definition 3.1 (Envy freedom). *An allocation \mathbf{x} with payments \mathbf{p} is envy free for valuation profile \mathbf{v} if no agent prefers the outcome of another agent to her own. Formally,*

$$\forall i, j, \quad v_i x_i - p_i \geq v_i x_j - p_j.$$

We first characterize envy-free outcomes in terms of the allocation. For a given allocation \mathbf{x} there may be several pricings \mathbf{p} for which the allocation is envy free. Since our objective is profit maximization we will characterize the \mathbf{p} corresponding to \mathbf{x} that gives the highest total revenue. The proof of this characterization as it is nearly identical to that of the analogous (and standard) characterization of incentive compatible mechanisms; we include it for completeness.

Definition 3.2. *An allocation is swap monotone if the allocation probabilities have the same order as the valuations of the agents, i.e., $x_i \geq x_{i+1}$ for all i . (Recall agents are ordered with $v_i \geq v_{i+1}$.)*




Figure 1: The solid curve depicts a swap monotone allocation as a function of the values (points). The shaded area corresponds to agent i 's payment p_i from Lemma 3.1.

Lemma 3.1. *In symmetric environments, an allocation \mathbf{x} admits a non-negative and individual rational payment \mathbf{p} such that (\mathbf{x}, \mathbf{p}) is envy free if and only if \mathbf{x} is swap monotone. If \mathbf{x} is swap monotone, then the maximum payments for which (\mathbf{x}, \mathbf{p}) is envy free satisfy, for all i , (Figure 1)*

$$\begin{aligned} p_i &= \sum_{j=i}^n (v_j - v_{j+1}) \cdot (x_i - x_{j+1}) \\ &= \sum_{j=i}^n v_j \cdot (x_j - x_{j+1}). \end{aligned}$$

Proof. Monotonicity and the payment identity imply envy-freeness: Suppose \mathbf{x} is swap monotone. Let \mathbf{p} be given as in the lemma. We verify that (\mathbf{x}, \mathbf{p}) is envy-free. There are two cases: if $i \leq j$, we have:

$$p_i - p_j = \sum_{k=i}^{j-1} v_k \cdot (x_k - x_{k+1}) \leq v_i \cdot \sum_{k=i}^{j-1} (x_k - x_{k+1}) = v_i \cdot (x_i - x_j),$$

and if $i \geq j$, we have:

$$p_i - p_j = - \sum_{k=j}^{i-1} v_k \cdot (x_k - x_{k+1}) \leq -v_i \cdot \sum_{k=j}^{i-1} (x_k - x_{k+1}) = v_i \cdot (x_i - x_j).$$

Rearranging the results of these calculations we have the definition of envy freedom.

Envy-freeness implies monotonicity: Suppose \mathbf{x} admits \mathbf{p} such that (\mathbf{x}, \mathbf{p}) is envy-free. By definition, $v_i x_i - p_i \geq v_i x_j - p_j$ and $v_j x_j - p_j \geq v_j x_i - p_i$. By summing these two inequalities and rearranging, $(x_i - x_j) \cdot (v_i - v_j) \geq 0$, and hence \mathbf{x} is swap monotone.

The maximum envy-free prices satisfy the payment identity: Agent i does not envy $i+1$ so $v_i x_i - p_i \geq v_i x_{i+1} - p_{i+1}$, or rearranging: $p_i \leq v_i(x_i - x_{i+1}) + p_{i+1}$. Given p_{i+1} the maximum p_i satisfies this inequality with equality. Letting $p_n = v_n x_n$ (the maximum individually rational payment) and induction gives the payment identity: $p_i = \sum_{j=i}^n v_j \cdot (x_j - x_{j+1})$. \square

Importantly, the above characterization leaves us free to speak of the (maximum) envy-free revenue of any swap monotone allocation \mathbf{x} on values \mathbf{v} , which we denote by $\text{EF}^{\mathbf{x}}(\mathbf{v})$. For any \mathbf{v} and any symmetric environment we will now solve for the envy-free optimal revenue, denoted by $\text{EFO}(\mathbf{v})$.




Figure 2: R and \bar{R} are the revenue curve and ironed revenue curve of the valuation profile $(6, 4, 4)$. The ironed virtual value of the high-value agent is 6, and the ironed virtual value of the two low-value agents are both $(12 - 6)/2 = 3$. E.g., the optimal EF revenue in the $k = 2$ unit environment is $\bar{R}(2) = 9$.

We will characterize the envy-free optimal revenue in terms of properties of the valuation profile \mathbf{v} . Given a valuation profile \mathbf{v} we denote the *revenue curve* by $R^{\mathbf{v}}(i) = i \cdot v_i$ for $i = \{1, \dots, n\}$ (recall v_i 's are indexed in decreasing order). For convenience we also let $R^{\mathbf{v}}(0) = R^{\mathbf{v}}(n+1) = 0$. The *ironed revenue curve*, denoted $\bar{R}^{\mathbf{v}}(i)$, is the minimum concave function that upper-bounds $R^{\mathbf{v}}$. Likewise, define the *virtual valuation function* $\Phi^{\mathbf{v}}(v) = R^{\mathbf{v}}(i) - R^{\mathbf{v}}(i-1)$ and the *ironed virtual valuation function* $\bar{\Phi}^{\mathbf{v}}(v) = \bar{R}^{\mathbf{v}}(i) - \bar{R}^{\mathbf{v}}(i-1)$, where $i \in \{1, \dots, n+1\}$ is such that $v \in [v_i, v_{i-1})$. (We set $v_0 = \infty$ for notational convenience.) See Figure 2.

$R^{\mathbf{v}}(i)$ is the best envy-free revenue one can get from serving exactly i agents at the same price deterministically. Consider a 2-unit auction example with one high-value agent with value 6 and two low-value agents with value 4. It is envy free to serve one high-value agent and one low-value agent at price 4, achieving revenue $R(2) = 8$. Interestingly, this is not optimal. The following allocation and payments are also envy-free: serve the high-value agent with probability 1 at price 5, and serve a low-value agent chosen at random at price 4. Both units are always sold and the total revenue is $\bar{R}(2) = 9$. In what follows we will derive that this revenue is optimal among all envy-free outcomes.

Lemma 3.2. *The (maximum) envy-free revenue of a swap monotone allocation \mathbf{x} satisfies:*

$$\text{EF}^{\mathbf{x}}(\mathbf{v}) = \sum_{i=1}^n R^{\mathbf{v}}(i) \cdot (x_i - x_{i+1}) = \sum_{i=1}^n \Phi^{\mathbf{v}}(v_i) \cdot x_i.$$

Proof. The proof is by the following equalities:

$$\begin{aligned} \text{EF}^{\mathbf{x}}(\mathbf{v}) &= \sum_{i=1}^n p_i = \sum_{i=1}^n \sum_{j=i}^n v_j \cdot (x_j - x_{j+1}) \\ &= \sum_{i=1}^n i v_i \cdot (x_i - x_{i+1}) = \sum_{i=1}^n R(i) \cdot (x_i - x_{i+1}) \\ &= \sum_{i=1}^n (R(i) - R(i-1)) \cdot x_i = \sum_{i=1}^n \Phi^{\mathbf{v}}(v_i) \cdot x_i. \end{aligned} \quad \square$$

An implication of the characterization of the envy-free revenue of a pricing as its *virtual surplus*, i.e., $\sum_i \Phi(v_i) x_i$, suggests that to maximize revenue, the allocation should maximize virtual surplus subject to swap monotonicity (and feasibility). In symmetric environments with monotone virtual valuation functions, the maximization of virtual surplus results in a swap monotone allocation. In

general symmetric environments, the allocation that maximizes *ironed* virtual surplus is both swap monotone and revenue optimal among all swap monotone allocations.

Lemma 3.3. *In a symmetric environments, the allocation that maximizes ironed virtual surplus with ties broken randomly is swap monotone.*

Proof. Suppose $\bar{\Phi}(v_i) > \bar{\Phi}(v_j)$ then $x_i \geq x_j$; otherwise, swapping x_i for x_j would have higher ironed virtual surplus. Suppose $\bar{\Phi}(v_i) = \bar{\Phi}(v_j)$, then $x_i = x_j$ because of random tie breaking and the symmetry of the environment. \square

Theorem 3.4. *In any symmetric environment with any valuation profile \mathbf{v} , the allocation \mathbf{x} that maximizes ironed virtual surplus w.r.t. $\bar{\Phi}^{\mathbf{v}}$ maximizes envy-free revenue among all swap-monotone allocations. I.e., $\text{EFO}(\mathbf{v}) = \text{EF}^{\mathbf{x}}(\mathbf{v})$.*

This theorem is proved by a useful lemma that relates revenue to ironed virtual surplus.

Lemma 3.5. *For any swap-monotone allocation \mathbf{x} on valuation profile \mathbf{v} ,*

$$\text{EF}^{\mathbf{x}}(\mathbf{v}) \leq \sum_{i=1}^n \bar{\Phi}^{\mathbf{v}}(v_i) \cdot x_i = \sum_{i=1}^n \bar{R}^{\mathbf{v}}(i) \cdot (x_i - x_{i+1}),$$

with equality holding if and only if $x_i = x_{i+1}$ whenever $\bar{R}^{\mathbf{v}}(i) > R^{\mathbf{v}}(i)$.

Proof. To show the inequality, we have:

$$\begin{aligned} \text{EF}^{\mathbf{x}}(\mathbf{v}) &= \sum_{i=1}^n R(i) \cdot (x_i - x_{i+1}) \\ &= \sum_{i=1}^n \bar{R}(i) \cdot (x_i - x_{i+1}) \\ &\quad - \sum_{i=1}^n (\bar{R}(i) - R(i)) \cdot (x_i - x_{i+1}) \\ &\leq \sum_{i=1}^n \bar{R}(i) \cdot (x_i - x_{i+1}), \end{aligned}$$

where we use the fact that $\bar{R}(i) \geq R(i)$ and $x_i \geq x_{i+1}$. Clearly the equality holds if and only if $x_i = x_{i+1}$ whenever $\bar{R}(i) > R(i)$. \square

Proof of Theorem 3.4. Consider \mathbf{x} that optimizes ironed virtual surplus with random tie breaking and also consider any other swap monotone \mathbf{x}' . Note that whenever $\bar{R}(i) > R(i)$, we have $\bar{\Phi}^{\mathbf{v}}(v_i) = \bar{\Phi}^{\mathbf{v}}(v_{i+1})$ for which random tie-breaking implies $x_i = x_{i+1}$. Therefore \mathbf{x} satisfies Lemma 3.5 with equality and it is optimal for the summation of the equality, whereas \mathbf{x}' satisfies it with inequality and may not be optimal for the summation. Thus $\text{EF}^{\mathbf{x}}(\mathbf{v}) \geq \text{EF}^{\mathbf{x}'}(\mathbf{v})$ and \mathbf{x} is revenue optimal. \square

As an example of this theorem, consider the position auction environment with weights $w_1 \geq w_2 \geq \dots \geq w_n$. An ironed virtual surplus maximizer assigns agents with higher ironed virtual values to slots with larger click probabilities, breaking ties randomly, ignoring agents with negative ironed virtual values. The ironed virtual surplus, and thus revenue, is $\sum_{\{i : \bar{\Phi}(v_i) \geq 0\}} \bar{\Phi}(v_i) \cdot w_i$, which can be read off the revenue curve, e.g., Figure 2.

Importantly, ironed virtual surplus maximization for position auctions is ordinal, i.e., only the order of the ironed virtual values matters. The optimal envy-free outcome can then rephrased as follows: First, tentatively assign the agents to slots in order of their values. Second, randomly

permute the order of each group of agents with equal ironed virtual surplus. In section 6 we will discuss consequences for environments for which surplus maximization is ordinal.

In downward-closed permutation environments, after the set system is realized, we find the allocation that maximizes ironed virtual surplus. With the appropriate payments, this outcome is envy free given the permutation. Importantly, this optimization is not ordinal.

4 Incentive Compatibility versus Envy Freedom

The major challenge in designing and analyzing incentive compatible mechanisms is that the incentive constraint binds across all possible misreports of the agents. We therefore view a mechanism as an allocation rule and payment rule pair where $\mathbf{x}(\mathbf{v})$ and $\mathbf{p}(\mathbf{v})$ denote the allocation and payments as a function of the agent values.

Definition 4.1 (Incentive Compatibility). *A mechanism is incentive compatible if no agent prefers the outcome when misreporting her value to the outcome when reporting the truth. Formally,*

$$\forall i, z, \mathbf{v}, \quad v_i x_i(\mathbf{v}) - p_i(\mathbf{v}) \geq v_i x_i(z, \mathbf{v}_{-i}) - p_i(z, \mathbf{v}_{-i}),$$

where (z, \mathbf{v}_{-i}) is obtained from \mathbf{v} with v_i replaced by z .

Definition 4.2 (Value Monotonicity). *An allocation rule is value monotone if the probability that an agent is served is monotone non-decreasing in her value, i.e., $x_i(z, \mathbf{v}_{-i})$ is non-decreasing in z for all agents i .*

The following well-known theorem characterizes ex post IC mechanisms.

Theorem 4.1. (Myerson, 1981) *An allocation rule $\mathbf{x}(\cdot)$ admits a non-negative and individually rational payment rule $\mathbf{p}(\cdot)$ such that (\mathbf{x}, \mathbf{p}) is incentive compatible if and only if $\mathbf{x}(\cdot)$ is value monotone, and the uniquely determined payment rule is:*

$$p_i(\mathbf{v}) = v_i x_i(\mathbf{v}) - \int_0^{v_i} x_i(z, \mathbf{v}_{-i}) dz.$$

Because the payments are uniquely determined by the allocation rule, for any allocation rule $\mathbf{x}(\cdot)$, we let $\text{IC}^{\mathbf{x}}(\mathbf{v})$ denote the IC revenue from running $\mathbf{x}(\cdot)$ over \mathbf{v} .

We now compare envy-free revenue to incentive-compatible revenue for ironed virtual surplus optimizers in permutation environments, i.e., where agents are assigned to roles in the set system via a random permutation. We show that these quantities are often within a factor of two of each other.

First we lower bound IC revenue by half of the maximum envy-free revenue under a technical condition. In the following we use $\text{IC}_i^{\bar{\Phi}}(\mathbf{v})$ and $\text{EF}_i^{\bar{\Phi}}(\mathbf{v})$ to denote the IC and EF revenue from agent i by applying the ironed virtual surplus maximizer $\bar{\Phi}$, respectively.

Lemma 4.2. *For downward-closed permutation environments, all valuations \mathbf{v} , and $\bar{\Phi}$, the ironed virtual valuation function corresponding to some \mathbf{v}' obtained from \mathbf{v} by setting a subset of agents' values to be 0, we have that $\text{IC}_i^{\bar{\Phi}}(\mathbf{v}) \geq \frac{1}{2} \text{EF}_i^{\bar{\Phi}}(\mathbf{v})$ for all i .*

Proof. Let $\mathbf{x}(\cdot)$ denote the allocation rule of the ironed virtual surplus optimizer $\bar{\Phi}$. By the assumption of the lemma, for all j , $\bar{\Phi}(z)$ is constant for all $z \in [v_{j+1}, v_j]$, and hence the IC allocation rule in fact maps each $z \in [v_{j+1}, v_j]$ to $x_i(v_{j+1}, \mathbf{v}_{-i})$.

By Lemma 4.1, $\text{IC}_i^{\bar{\Phi}}(\mathbf{v})$ is equal to $\sum_{j=i}^n (v_j - v_{j+1}) \cdot (x_i(\mathbf{v}) - x_i(v_{j+1}, \mathbf{v}_{-i}))$ which, referring to Figure 3, equals the area above the IC curve and below the horizontal dotted line. On the other hand, $\text{EF}_i^{\bar{\Phi}}(\mathbf{v})$ is equal to $\sum_{j=i}^n (v_j - v_{j+1}) \cdot (x_i(\mathbf{v}) - x_{j+1}(\mathbf{v}))$, which similarly corresponds to the area above the EF curve and below the horizontal dotted line. It suffices to prove that: $x_i(\mathbf{v}) - x_i(v_{j+1}, \mathbf{v}_{-i}) \geq \frac{1}{2} \cdot (x_i(\mathbf{v}) - x_{j+1}(\mathbf{v}))$. Note that $x_i(v_{j+1}, \mathbf{v}_{-i}) = x_{j+1}(v_{j+1}, \mathbf{v}_{-i})$ as now agents i and $j+1$ have the same value, this is equivalent to $x_i(\mathbf{v}) + x_{j+1}(\mathbf{v}) \geq x_i(v_{j+1}, \mathbf{v}_{-i}) + x_{j+1}(v_{j+1}, \mathbf{v}_{-i})$.

The last inequality says that the total winning probability of agent i and $j+1$ can only decrease if agent i lowers her bid to v_{j+1} . To prove this, we fix the permutation that maps agents to roles of the set system, and show that the number of winning agents from i and $j+1$ can only be lower after agent i decreases her value. There are two cases to verify: (1) Agent i wins after the decrease. Then before the decrease, agent i had higher value, and the optimal feasible set would be the same. (2) Agent $j+1$ wins and agent i loses after the decrease. Then before the decrease, at least one of agents i and $j+1$ would win. \square




Figure 3: Depiction of EF allocation and IC allocation rule from which the payments for agent i are computed. The EF allocation curve maps each value in $[v_{j+1}, v_j]$ to $x_{j+1}(\mathbf{v})$, and the IC allocation curve maps each z to $x_i(z, \mathbf{v}_{-i})$.

In matroid permutation environments, envy-free revenue upper-bounds incentive-compatible revenue.

Lemma 4.3. *For matroid permutation environments, all valuations \mathbf{v} , and all ironed virtual valuation functions $\bar{\Phi}$, for all agent i , $\text{EF}_i^{\bar{\Phi}}(\mathbf{v}) \geq \text{IC}_i^{\bar{\Phi}}(\mathbf{v})$.*

Proof. Recall that $\text{EF}_i^{\bar{\Phi}}(\mathbf{v}) = \sum_{j=i}^n (v_j - v_{j+1}) \cdot (x_i(\mathbf{v}) - x_{j+1}(\mathbf{v}))$ and $\text{IC}_i^{\bar{\Phi}}(\mathbf{v}) = \int_0^{v_i} (x_i(\mathbf{v}) - x_i(z, \mathbf{v}_{-i})) dz$. By the monotonicity of $x_i(z, \mathbf{v}_{-i})$ in z , $\text{IC}_i^{\bar{\Phi}}(\mathbf{v})$ is upper-bounded by $\sum_{j=i}^n (v_j - v_{j+1}) \cdot (x_i(\mathbf{v}) - x_i(v_{j+1}, \mathbf{v}_{-i}))$. Recall that $x_i(v_{j+1}, \mathbf{v}_{-i}) = x_{j+1}(v_{j+1}, \mathbf{v}_{-i})$. It suffices to prove that $x_{j+1}(\mathbf{v}) \leq x_{j+1}(v_{j+1}, \mathbf{v}_{-i})$. To see this, ironed virtual surplus maximizers are greedy algorithms in matroid permutation settings, and if agent i decreases her bid to v_{j+1} , agent $j+1$ is less likely to be blocked by i who was earlier in the greedy order, and is hence more likely to be allocated. \square

There are downward-closed permutation environments where the envy-free optimal revenue does not upper-bound the incentive-compatible revenue of all virtual surplus maximizers. The proof is by example and can be found in Appendix A.2; for this example the amount that the IC revenue exceeds the EF revenue is a small fraction of the total EF revenue.

Lemma 4.4. *There exists a downward-closed permutation environment and valuation profile \mathbf{v} , such that if $\bar{\Phi} = \bar{\Phi}^{\mathbf{v}}$ is the ironed virtual valuation function of \mathbf{v} , then $\text{IC}^{\bar{\Phi}}(\mathbf{v}) > \text{EF}^{\bar{\Phi}}(\mathbf{v})$.*

5 Prior-free Mechanism Design and Benchmarks

As discussed previously, no incentive-compatible mechanism obtains an optimal profit point-wise on all possible valuation profiles. Therefore, to obtain point-wise guarantees, the literature on prior-free mechanism design looks for the incentive compatible mechanism that minimizes, over valuation profiles, its worst-case ratio to a given performance benchmark. It is important to identify a good benchmark for such an analysis to be meaningful.

If the designer had a prior distribution over the agent valuations then she could design the mechanism that maximizes revenue in expectation over this distribution. This is the approach of Bayesian optimal mechanism design as characterized by Myerson (1981) and refined by Bulow and Roberts (1989). Given a distribution F , virtual values and revenue curves can be derived. The optimal mechanism is the one that maximizes ironed virtual surplus.

Theorem 5.1. *(Myerson, 1981) When values are i.i.d. from distribution F the optimal mechanism, ICO^F , is the ironed virtual surplus optimizer for $\bar{\Phi}$ corresponding to F .*

If the agent values are indeed drawn from a prior distribution, but the designer is unaware of the distribution, then a reasonable objective might be to design a mechanism that is a good approximation to the optimal mechanism for any unknown distribution. This *prior-independent* objective is a relaxation of our prior-free objective.

One important criterion for a prior-free benchmark is that its approximation should imply prior-independent approximation: if a mechanism is a constant approximation to the benchmark, then for a relevant class of distributions, it should be a constant approximation to the Bayesian optimal mechanism under any distribution from the class.

For matroid permutation environments, Lemma 4.3 implies that for any values \mathbf{v} the optimal envy-free revenue $\text{EFO}(\mathbf{v})$ (which is at least the envy-free revenue of any ironed virtual surplus optimizer) is at least the incentive compatible revenue of any ironed virtual surplus optimizer. By Theorem 5.1, the Bayesian optimal mechanism is an ironed virtual surplus optimizer so $\text{EFO}(\mathbf{v})$ upper-bounds its revenue. Consequently, a prior-free β -approximation to EFO is also a prior-independent β -approximation for all distributions.

Unfortunately, even for simple the digital good environment it is not possible to obtain a prior-free constant approximation to EFO (see Goldberg et al., 2006b). This impossibility arises because it is not possible to approximate the highest value v_1 . For essentially the same reason, it is not possible to design a prior-independent constant approximation for all distributions. We therefore restrict attention to the large family of distributions with tails that are not too irregular.

Definition 5.1 (Tail Regularity). *A distribution F is n -tail regular if in n -agent 1-unit environments, the expected revenue of the Vickrey auction is a 2-approximation to that of the Bayesian optimal mechanism.*

The definition of tail regularity is implied by Myerson's regularity assumption via the main theorem of Bulow and Klemperer (1996). The intuition for the definition is the following. For n -agent 1-unit environments, all the action happens in the tail of the distribution, i.e, values v for which $F(v) \approx 1 - 1/n$; therefore, irregularity of the rest of the distribution does not have much

consequence on revenue. Tail regularity, then, restates the Bulow-Klemperer consequence, as a constraint on the tail of the distribution and leaves the rest unconstrained.

We now define the benchmark for prior-free mechanism design. Approximation of this benchmark guarantees prior-independent approximation of all n -tail-regular distributions.

Definition 5.2. *The envy-free benchmark is $\text{EFO}^{(2)}(\mathbf{v}) = \text{EFO}(\mathbf{v}^{(2)})$ where $\mathbf{v}^{(2)} = (v_2, v_2, v_3, \dots, v_n)$.*

Theorem 5.2. *With any n -agent matroid permutation environment, any n -tail-regular distribution F , and any β -approximation mechanism to $\text{EFO}^{(2)}$, the expected revenue of the mechanism with valuations \mathbf{v} drawn i.i.d. from F is a 3β -approximation to the optimal mechanism for F .*

Proof sketch. (The full proof in Appendix A.3.) We focus on showing the result for k -unit auctions; the reduction in Section 6 will enable us to easily generalize this to matroid permutation environments. We use tail regularity to get a bound on the payment from the highest agent in terms of the Vickrey auction revenue, v_2 . Of course, $\text{EFO}^{(2)}(\mathbf{v})$ is at least v_2 , and so the payment from the highest agent is at most $2 \text{EFO}^{(2)}(\mathbf{v})$ (in expectation over i.i.d. draws of \mathbf{v} from F). The second part of the argument involves bounding the total payments of agents $\{2, \dots, n\}$, point-wise from above, by $\text{EFO}^{(2)}(\mathbf{v})$. This is possible by Lemma 4.3 and detailed analysis of k -unit auction payments. \square

It is useful to compare the EFO benchmark to ones proposed in the literature that are based on the Vickrey-Clarke-Groves (VCG) mechanism with the best (for the particular valuation profile \mathbf{v}) reserve price (e.g., Goldberg et al., 2006a; Hartline and Roughgarden, 2009). The VCG mechanism with a reserve price first rejects all agents whose values do not meet the reserve, then it serves the remaining agents to maximize the surplus (sum of values).

The VCG-with-reserve benchmark can be expressed as an ironed virtual surplus optimizer, and so by Lemma 4.3, in matroid permutation environments, EFO is no smaller. For a digital good, EFO and VCG-with-reserve are identical. For multi-unit auctions EFO is at most twice VCG-with-reserve. For matroid permutation environments EFO can be (almost) a logarithmic factor larger than VCG-with-reserve. Therefore, for general environments the EFO-based benchmark results in stronger approximation guarantees.

Theorem 5.3. *For any multi-unit environment and valuation profile \mathbf{v} , $\text{EFO}(\mathbf{v}) \leq 2 \text{VCG-with-reserve}(\mathbf{v})$; furthermore, for k units there exists a valuation profile \mathbf{v} such that $\text{EFO}(\mathbf{v}) \geq (2 - \frac{1}{k}) \text{VCG-with-reserve}(\mathbf{v})$.*

Proof. In the k -unit environment, the VCG-with-reserve benchmark on \mathbf{v} is simply $\max_{i \leq k} R(i)$. The envy-free benchmark is $\text{EFO}(\mathbf{v}) = \max_{i \leq k} \bar{R}(i)$. Recall that \bar{R} is concave. If \bar{R} attains its maximum at $r < k$ then $R(r) = \bar{R}(r)$ and both benchmarks obtain the same revenue and the first part of the lemma holds; so assume that \bar{R} attains its maximum on $\{1, \dots, k\}$ at k .

Suppose R is ironed on interval $i < k < j$. Then $\bar{R}(k)$ is the convex combination of $R(i)$ and $R(j)$ as

$$\begin{aligned} \bar{R}(k) &= \frac{j-k}{j-i} R(i) + \frac{k-i}{j-i} R(j) \\ &\leq R(i) + \frac{k}{j} R(j) \leq R(i) + R(k) \leq 2 \text{VCG-with-reserve}(\mathbf{v}). \end{aligned} \tag{1}$$

The second inequality follows because $v_k \geq v_j$ implies that $kv_k \geq \frac{k}{j}jv_j$; the final equality follows because both $R(i)$ and $R(k)$ are feasible revenues for VCG with reserve.

The above bound is almost tight. Let $v_1 = k$ and $v_2 = v_3 = \dots = v_n = 1$. The VCG-with-reserve benchmark obtains revenue k whereas $\bar{R}(k) = \frac{n-k}{n-1}k + \frac{k-1}{n-1}n$, from equation (1), which approaches $2k - 1$ as $n \rightarrow \infty$. \square

Lemma 5.4. *There exists a distribution F and n -agent matroid environment for which VCG with any reserve price is an $\Omega(\log n / \log \log n)$ -approximation to the Bayesian optimal mechanism for F .*

Proof sketch. (A full proof is in Appendix A.1) We construct a set system and an irregular distribution with a jagged revenue curve that has $m = \Omega(\frac{\log n}{\log \log n})$ deep “trenches”, such that the following are true. (1) the Bayesian optimal mechanism for F gets about $1/m$ fraction of its revenue from each trench via ironing. (2) VCG with reserve gets similar amount of good revenue from at most one of the trenches by setting an appropriate reserve price, but only gets low revenue from the other trenches due to the lack of ironing. In total, VCG with reserve only gets about $1/m$ fraction of the Bayesian optimal revenue. \square

We conclude this section with a summary. For multi-unit environments we have given Bayesian justification for the envy-free benchmark EFO⁽²⁾. A mechanism that approximates this prior-free benchmark simultaneously approximates the Bayesian optimal mechanisms for most i.i.d. distributions (i.e., those satisfying the relatively unrestrictive tail-regularity assumption). In the next section we will give a reduction from matroid permutation and position environments to multi-unit environments. This reduction implies that the envy-free benchmark is also Bayesian justified for these environments. While we have not given formal justification for the envy-free benchmark in downward-closed permutation environments (because of Lemma 4.4); we believe its approximation is still of interest.

6 Multi-unit, Position, and Matroid Permutation Environments

In this section we consider matroid permutation, position auction, and multi-unit environments. We show that for both incentive-compatible mechanism design and envy-free outcomes, these environments are closely related. In fact, for either IC or EF, respectively, the optimal mechanisms across these environments are the same and approximation mechanisms give the same approximation factor. As an example, we will focus on approximating the envy-free benchmark EFO⁽²⁾ (Definition 5.2) with a prior-free mechanism. Our solution will be via a two-step reduction: we reduce matroid permutation to position auction environments, which we then reduce to multi-unit environments.

Recall that in a multi-unit environment it is feasible to serve any set of agents of cardinality at most some given k . In position auction environments there are weights $w_1 \geq w_2 \geq \dots \geq w_n$ for positions and feasible outcomes are partial assignments of agents to positions. In matroid permutation environments there is a feasibility constraint given by independent sets of a matroid, but the roles of the agents are assigned by random permutation.

The property of these three environments that enables this reduction is that in each environment the greedy algorithm on ironed virtual values (with ties broken randomly) obtains the maximum ironed virtual surplus. The greedy algorithm works as follows: order the agents by ironed virtual value and serve each agent in this order if her ironed virtual value is positive and if doing so is feasible given the set of agents previously served. Notice that the only information needed to perform this maximization is the ordering on the agents’ ironed virtual values (but not their magnitudes).

Definition 6.1. *The characteristic weights $w_1 \geq w_2 \geq \dots \geq w_n$ of a matroid environment are as follows: choose any valuation profile \mathbf{v} with all distinct values, assign the agents to elements in the matroid via a random permutation, run the greedy algorithm w.r.t. \mathbf{v} , and define w_i to be the probability that i th largest valued agent is served.*

6.1 Reduction for Ironed Virtual Surplus Maximizers

We first show ironed virtual surplus optimization in the three environments is equivalent.

Proposition 6.1. *The ironed virtual surplus maximizing assignment (and its virtual surplus) is equal in expectation in the following environments:*

1. *a matroid permutation environment with characteristic weights \mathbf{w} ,*
2. *a position auction environment with weights \mathbf{w} ,*
3. *a convex combination of multi-unit environments where k units are available with probability $w_k - w_{k+1}$ for $k \in \{1, \dots, n\}$ and $w_{n+1} = 0$.*

Proof. Fix a tie-breaking rule, which induces an ordering on the agents. Consider the greedy algorithm on the agents with non-negative $\bar{\Phi}$ values according to this ordering. The j -th agent with non-negative $\bar{\Phi}$ value in this ordering (1) gets allocated with probability w_j in the matroid permutation setting by definition of characteristic weights, (2) gets assigned to position j in the position auction and hence gets allocated with probability w_j , and, (3) gets allocated in k -unit auction for each $k \geq j$, and hence has probability $\sum_{k \geq j} (w_k - w_{k+1}) = w_j$ of being served in the convex combination setting. Taking expectation over all tie-breaking orders, agent i has the same probability of being served in the three settings. \square

The following corollary is immediate.

Corollary 6.2. *For any valuation profile \mathbf{v} and any weights \mathbf{w} , the envy-free optimal revenue is the same in each of the environments of Proposition 6.1.*

A basic fact about incentive compatibility is that it is closed under convex combination, i.e., a randomization over two incentive compatible mechanisms is incentive compatible: truthtelling is an optimal strategy in each, and so it remains an optimal strategy in the combination.

We now illustrate how to use Proposition 6.1 to show that an incentive compatible prior-free approximation mechanisms for multi-unit environments can be adapted to give the same approximation factor in position auction and matroid permutation environments. Consider the following incentive compatible mechanism.

Definition 6.2. *The Random Sampling Empirical Myerson (RSEM) mechanism does the following: (discussion of payments omitted)*

1. *randomly partition the population of agents $N = \{1, \dots, n\}$ into two sets by flipping a fair coin for each agent,*
2. *designate the set containing the highest-valued agent as the market M and the other set as the sample S ,*
3. *calculate the ironed virtual surplus function $\bar{\Phi}^S$ for the sample S , and,*

4. serve a feasible subset of M to maximize ironed virtual surplus with respect to $\bar{\Phi}^S$ and reject all other agents.

Lemma 6.3. *In any single-dimensional agent environment, RSEM is incentive compatible.*

Proof. Notice that RSEM is monotone: An agent in S loses unless she raises her bid to beat the highest-valued agent (in which case the roles of S and M are reversed). An agent in M wins when the virtual surplus maximizing set contains the agent. If she raises her bid, she (weakly) increases her virtual value thus increasing the virtual surplus of any set containing her, while the virtual surplus of other sets remain the same. Therefore, she continues to win. By Theorem 4.1 monotonicity implies that, with the appropriate payments, RSEM is incentive compatible. \square

The proof of the following theorem is technical and we defer discussion of it to Section 6.3.

Theorem 6.4. *In multi-unit environments, RSEM is a 12.5-approximation to the envy-free benchmark $\text{EFO}^{(2)}(\mathbf{v})$.*

Notice that this mechanism can easily be generalized to other downward-closed environments. It remains incentive compatible for these environments because it is essentially an ironed virtual surplus optimizer on the set M , and furthermore, it is incentive compatible even if the permutation that assigns agents to the set system is fixed. As a final corollary of Proposition 6.1, we can view RSEM's revenue in the matroid permutation or position auction environment as the analogous convex combination of its revenue in multi-unit auction environments.

Corollary 6.5. *In matroid permutation and position environments, RSEM is a prior-free 12.5-approximation to the envy-free revenue $\text{EFO}^{(2)}(\mathbf{v})$.*

6.2 General Reduction

The following prior-free approximations are essentially the best known for digital-good and multi-unit environments. Notably, the mechanism from Corollary 6.8 below, is not based on ironed virtual surplus maximization and therefore Proposition 6.1 cannot be applied to a construct matroid permutation or position auction mechanism from it.

Lemma 6.6. *(Ichiba and Iwama, 2010) In the digital good environment, there is a prior-free incentive compatible 3.12-approximation to $\text{EFO}^{(2)}(\mathbf{v})$.*

We now give an approximate reduction from multi-unit auctions to digital good auctions. This construction and the proof that the resulting mechanisms incentive compatibility are standard. See, e.g., Myerson (1981), Goldberg et al. (2006a), and Aggarwal and Hartline (2006).

Definition 6.3 (Multi-unit Reduction). *Given a k -agent digital good auction, we construct the following k -unit auction:*

1. Simulate the k -unit Vickrey auction.
2. Simulate the k -agent digital good auction on the k winners of the Vickrey auction.
3. Serve the agents who win in both stages and charge them the maximum of their simulation payments; reject all other agents.

Theorem 6.7. *Given any digital-good auction that β -approximates to the envy-free benchmark (resp. Bayesian optimal auction), the multi-unit auction from the reduction is a 2β -approximation to the envy-free benchmark (resp. Bayesian optimal auction).*

Proof. The digital good auction is a β -approximation to the envy-free benchmark on the top k agents. The envy-free benchmark on the top k agents is equal to the VCG-with-reserve benchmark for the full set of agents (both are equal to $\max_{i \leq k} R(i)$). Theorem 5.3 states that the VCG-with-reserve benchmark is a 2-approximation to the envy-free benchmark. Therefore, the multi-unit auction from the reduction is a 2β -approximation to the envy-free benchmark. \square

Corollary 6.8. *In multi-unit environments, there is an incentive compatible prior-free 6.24-approximation to $\text{EFO}^{(2)}(\mathbf{v})$.*

We now show how to construct, from any multi-unit auction, a position auction and matroid permutation mechanism that has the exact same outcome (in expectation) as a convex combination of multi-unit auctions (as in Proposition 6.1). The challenge here is the distinct interfaces to the environment: in multi-unit auctions we are given a supply constraint k and we need to specify a set of at most k winners, whereas in position auctions, we are given weights and need to output a partial assignment of agents to positions.

Definition 6.4 (Position Auction Reduction). *Given k -unit auction mechanisms for $k \in \{1, \dots, n\}$, we construct the following mechanism for the position auction environment with weights \mathbf{w} :*

1. *Introduce n dummy agents and n dummy positions into the system, indexed by $\{n+1, \dots, 2n\}$. Correspondingly, we pad weights \mathbf{w} and valuation profile \mathbf{v} with zeros such that they have dimension $2n$.*
2. *For each $k \in \{1, \dots, n\}$, simulate the k -unit auction on valuation profile \mathbf{v} , and give the unallocated leftover units to the dummy agents arbitrarily for free. Let the resulting allocation of all $2n$ agents be $\mathbf{x}^{(k)}$.*
3. *Calculate the probability that each agent is served in the convex combination: $x_i = \sum_{k=1}^n x_i^{(k)} (w_k - w_{k+1})$, for $i \in \{1, \dots, 2n\}$.*
4. *Solve for a set of permutation matrices $P_t \in \{0, 1\}^{2n \times 2n}$ and nonnegative weights r_t with $\sum_t r_t = 1$ such that $\sum_t r_t \cdot P_t \cdot \mathbf{w} = \mathbf{x}$.*
5. *With probability r_t , assign agents to positions according to the permutation specified by P_t .*
6. *Discard dummy agents and dummy position assignments.*

To justify step 4, one can verify that \mathbf{w} majorizes \mathbf{x} in the sense that $\sum_{i=1}^k w_i \geq \sum_{i=1}^k x_i$ for $k \in \{1, \dots, 2n\}$, with equality holding for $k = 2n$. Therefore by a theorem of Rado (1952), the desired permutation matrices and weights exist. The following consequences are immediate.

Lemma 6.9. *The resulting mechanism for position auction with weights \mathbf{w} obtained from the above reduction has the same allocation as the convex combination of k -unit auctions with $(w_k - w_{k+1})$'s as probabilities.*

Lemma 6.10. *Given an incentive compatible multi-unit auction, the mechanism from the position auction reduction is also incentive compatible.*

Definition 6.5 (Matroid Permutation Reduction). *Given a position auction mechanism for weights \mathbf{w} , we construct the following mechanism for matroid permutation environment with characteristic weights \mathbf{w} :*

1. *We run the position auction and for $i = 1, \dots, n$, let j_i be the position assigned to agent i , or $j_i = \perp$ if i is not assigned a position.*
2. *Reject all agents i with $j_i = \perp$.*
3. *Run the greedy algorithm in the matroid permutation environment with agent i 's value reset to j_i .*

The following conclusions are immediate.

Lemma 6.11. *The resulting mechanism for matroid permutation environment obtained from the above reduction has the same allocation as the position auction.*

Lemma 6.12. *Given an incentive-compatible position auction, the mechanism from the matroid permutation reduction is incentive compatible (in matroid permutation environments).*

Theorem 6.13. *The factor β to which there is a prior-free incentive-compatible approximation of $\text{EFO}^{(2)}(\mathbf{v})$ is the same for multi-unit, position, and matroid permutation environments.*

Corollary 6.14. *There is a prior-free incentive-compatible 6.24-approximation to $\text{EFO}^{(2)}(\mathbf{v})$ in matroid permutation and position environments.*

There are two weakness in the reductions implied by Theorem 6.13 in comparison to those implied by Proposition 6.1. Recall that for the latter, ironed virtual surplus maximizations are via the greedy algorithm, and so the reductions were algorithmically trivial. In contrast, Theorem 6.13 requires knowledge of the characteristic weights to run the construction, these weights may be hard to compute. In addition the mechanism that results from the matroid permutation reduction is only incentive compatible if the agents are assigned to roles in the matroid via a random permutation as suggested in the model. In contrast, RSEM in matroid environments is incentive compatible without any random permutation (Lemma 6.3).

6.3 Multi-unit Analysis of RSEM

In this section we prove Theorem 6.4 which shows that in multi-unit environments RSEM (Definition 6.2) is a 12.5-approximation to the envy-free benchmark. We further assume that the number of units available is $k \geq 2$. This is without loss of generality because with $k = 1$ RSEM produces the same outcome as the Vickrey auction which is optimal with respect to the envy-free benchmark $\text{EFO}^{(2)}(\mathbf{v}) = v_2$.

The approximation factor is the product of three terms: $2 \times \frac{4}{3} \times 4.68 \approx 12.5$. Roughly, these three terms come from the following steps:

1. We bound the performance relative to the VCG-with-reserve benchmark (i.e., $\max_{i \leq k} R(i)$). By Theorem 5.3 this benchmark can be at most a factor of two below the envy-free benchmark.
2. We use an intermediary envy-free-like revenue which is (usually) at most $\frac{4}{3}$ of RSEM's incentive-compatible revenue.

3. Finally we employ a result due to Alaei et al. (2009) to show that the expected relative imbalance between the sample and the market is 4.68; in particular our intermediary envy-free-like revenue is a 4.68 fraction of the VCG-with-reserve benchmark in expectation.

To lay out the quantities we will be discussing, recall that RSEM partitions the agents into a sample S and market M where the highest valued agent, 1, is conditioned to be in M . The (ironed) virtual surplus for $\bar{\Phi}^S$ on M is then optimized by selling the k units to the agents with the k highest (positive) virtual values. Conditioned on the partitioning, let p and q denote the supremum and infimum values that correspond to the virtual value of the $k+1$ st highest agent in M . Notice that if an agent in M bids above p she always wins, if she bids below q she always loses, and if she bids on (q, p) she wins with some probability. Let i and j be the number of agents in the market that are above p and q , respectively, and observe that $i \leq k < j$. This sort of outcome was termed by Hartline and Roughgarden (2009) as a p - q -lottery.

Notice that if we had enough units it would be envy free (for M) to sell to the top j agents at price q for revenue $L(j) = qj$ or to sell to the top i agents at price p for revenue $L(i) = pi$. Envy-free revenue is linear and therefore with only k items we can linearly interpolate between these two revenues to obtain revenue (cf. Section 5, equation (1)):

$$L(k) = \frac{j-k}{j-i}L(i) + \frac{k-i}{j-i}L(j).$$

This is the envy-free intermediary we referred to in Step 2, above.²

The IC revenue of this p - q -lottery, and consequently of RSEM, is lower than the envy-free revenue given by $L(k)$. The top i agents each receive a unit and pay $p - (p - q)\frac{(k-i+1)}{j-i+1}$ (by Theorem 4.1). The next $j - i$ agents each receive a unit with probability $\frac{k-i}{j-i}$ and pay q . The revenue of RSEM for selling k units to M is thus,

$$\text{RSEM}(k) = (k - i)q + i \left(p - (p - q)\frac{(k-i+1)}{j-i+1} \right).$$

This is also a linear function of k and, while we require by assumption that $i \leq k < j$, formulaically $\text{RSEM}(j) = L(j)$.

We now derive a linear function that upper bounds $\bar{R}^S(k)$, the ironed revenue curve for the sample. This upper bound will be parameterized by the degree of imbalance between the sample and market.

Definition 6.6. *The imbalance λ of a partitioning (S, M) of agents $N = \{1, \dots, n\}$ is $\lambda = \max_{i \in N} \frac{Y_i}{i - Y_i}$ where $Y_i = |\{1, 2, \dots, i\} \cap S|$ is the number of the highest i agents in S .*

For a given imbalance λ the number of agents in S with value at least p (resp. q) is at most $i\lambda$ (resp. $j\lambda$). Define $H(j) = qj\lambda$, $H(i) = \min(pi\lambda, H(j))$, and $H(k)$ as the linear interpolation between points $(i, H(i))$ and $(j, H(j))$. The reason we take the minimum in the second step is that since $\bar{R}^S(k)$ is monotone, an upper bound on $\bar{R}^S(j)$ is an upper bound for all $k \in [i, j]$. For this definition, $H(k) \geq \bar{R}^S(k)$; we defer proof of this lemma for later.

Lemma 6.15. *For $H(k)$ defined above, $H(k)$ is at least the envy-free revenue of the sample, $\bar{R}^S(k)$.*

²This envy-free revenue is not the revenue-maximal one specified by the payment identity of Lemma 3.1.

Now we have three linear functions defined on $[i, j]$: $H(\cdot)$, the upper bound on the envy-free revenue from the sample; $L(\cdot)$, the envy-free intermediary for the market; and $\text{RSEM}(\cdot)$, the incentive compatible revenue of RSEM on the market. Simply, when $L(\cdot)$ is increasing (so $H(\cdot)$ is parallel to it), the maximal ratio between RSEM and H occurs at $k = i$. The following lemma shows that this ratio is $\frac{4}{3}\lambda$ which follows from the ratio of $\text{RSEM}(i)$ to $L(i)$ being $\frac{4}{3}$; the proof is deferred to later.

Lemma 6.16. *With imbalance λ and increasing $L(\cdot)$, the ratio of RSEM to H is at most $\frac{4}{3}\lambda$.*

When $L(\cdot)$ is decreasing then $H(\cdot)$ is constant and the ratio between RSEM and H is maximized at $k \in \{i, j\}$. Because the ratio of $H(j)$ to $L(j) = \text{RSEM}(j)$ is λ by definition; we must only show that the $\text{RSEM}(i) \geq \frac{3}{4}L(j)$. We do this in the following lemma, proof of which is deferred to later.

Lemma 6.17. *With imbalance λ and decreasing $L(\cdot)$, the ratio of RSEM to H is at most $\frac{4}{3}\lambda$.*

We conclude by decomposing Step 3 into two parts. The first is the ratio between the VCG-with-reserve benchmark on all agents N and the envy-free benchmark on the sample S . The second is the ratio between our upper bound $H(\cdot)$ on the envy-free revenue of the sample S to the revenue of $\text{RSEM}(\cdot)$. The latter is, by Lemmas 6.16 and 6.17, at most $\frac{4}{3}\lambda$. To calculate the former, the VCG-with-reserve benchmark sells to $r \leq k$ agents at price v_r for a total revenue of $R(r) = rv_r$. Of course, on the sample S it is envy-free to post price v_r as well; such a price is accepted by Y_r (from the definition of imbalance, Definition 6.6) agents for a total revenue of $v_r Y_r$. The optimal envy-free revenue for the sample is no smaller. We conclude:

Lemma 6.18. *The envy-free revenue of the sample satisfies $\bar{R}^S(k) \geq \frac{Y_r}{r} \cdot \text{VCG-with-reserve}(k)$.*

Combining the ratios from these two parts, $\frac{r}{Y_r}$ and $\frac{4}{3}\lambda$ we conclude that our revenue is governed by the expectation of random variable $X = \frac{r}{Y_r}\lambda$. Its expectation is bounded by the following lemma:

Lemma 6.19. *(Alaei et al., 2009) For all positive integers r , random variable $X = \frac{r}{Y_r}\lambda$, for Y_i and imbalance λ as defined in Definition 6.6 satisfies:*

$$\mathbf{E}[1/X] \geq 1/4.68.$$

Lemmas 6.15–6.19 and Theorem 5.3 combine to prove Theorem 6.4. We conclude with the proofs of the three deferred lemmas.

Proof of Lemma 6.15. Assume $H(k)$ is increasing. We show that $H(k)$ is at least the envy-free revenue of the sample, $\bar{R}^S(k)$. \bar{R}^S is ironed between values p and q ; therefore, it has a line segment that connects two points on the lines from the origin with slope p and q . These respective points are closer to the origin than $(i\lambda, pi\lambda)$ and $(j\lambda, qj\lambda)$, by the definition of the imbalance parameter λ . Therefore, these latter points are strictly above the revenue curve $\bar{R}^S(\cdot)$ and furthermore so is the line segment that connects them. Since the revenue curve is monotone we can shift these points to the left and they (and the line segment between them) remain above the revenue curve. The line $H(\cdot)$ is constructed in this manner. If, on the other hand, $H(k)$ not increasing (then by definition it is constant) then all we need is the end point $(j, H(j))$ to be above the ironed revenue curve for S which is implicit in the argument above. \square

Proof of Lemma 6.16. We show that if $L(\cdot)$ is increasing then $RSEM(i) \geq \frac{3}{4}L(i)$. We assume that $i \geq 2$ (otherwise, a similar analysis for $i = 1$ and $k = 2$ can be employed). To do this we give an upper bound on

$$\frac{L(i) - RSEM(i)}{L(i)} = \frac{p - q}{p(j - i + 1)} \leq \frac{1}{6}.$$

The lemma follows easily from this upper bound. (In fact this gives a bound of $\frac{5}{6}$; the ratio $\frac{3}{4}$ is tight for a different case.) Consider maximizing the above quantity subject to the constraints

- $2 \leq i < j$. (i and j are integers.)
- $p > q > 0$.
- $pi \leq qj$.

Rewriting the last of these as $q/p \geq i/j$, we get that

$$\left(1 - \frac{q}{p}\right) \frac{1}{j - i + 1} \leq \left(1 - \frac{i}{j}\right) \frac{1}{j - i + 1} = \frac{j - i}{j(j - i + 1)}.$$

For any fixed j , the above quantity is maximized at the smallest i , which is 2. So we need to maximize

$$\frac{j - 2}{j(j - 1)}.$$

over all integers $j \geq 3$. It is easy to see that it is maximized at $j = 3$ with value $1/6$. \square

Proof of Lemma 6.17. When $L(\cdot)$ is decreasing we show that $RSEM(i) \geq \frac{3}{4}L(j)$. Suppose that actually $L(i) > (j + 1)q$. (This is slightly stronger than the hypothesis in the lemma, since $L(j) = jq$.) RSEM as a function of k is linear and has a negative slope, so it decreases in the interval $[i, j]$. If the linear function were to be extended to j , then it would be exactly $L(j)$. Thus $RSEM(i) \geq L(j)$. (There is no loss of a factor.)

The other case is where $L(j) < L(i) < (j + 1)q$. In this case one can compare $RSEM(i)$ with $L(j)$ and show that they are within a factor of $(j + 1)/j$ which is at most $4/3$, since $j \geq 3$. \square

7 Downward-closed Permutation Environments

In this section, we will show that a variant of RSEM (recall Definition 6.2) approximates the envy-free benchmark by a constant factor.

Definition 7.1 (RSEM'). *The variant RSEM' is identical to RSEM except Step 4:*

- 4'. *find the feasible subset W of N (the full set of agents) to maximize ironed virtual surplus with respect to Φ^S , serve agents in $M \cap W$ (the winners from the market M) only, and reject all other agents.*

The proof we give that RSEM' is a good approximation to the envy-free benchmark is based on the fact that with large probability the sample and market satisfy a natural balanced condition. This condition requires that, for all prefixes of the agents sorted by value, a good fraction of these agents are in each of the market and sample. The proof then has three main steps: show the probability of balance is high, show balance implies that the IC revenue of RSEM' (on the market M) is close to the optimal EF revenue for the sample S , and show the expected optimal EF revenue of the sample is close to the envy-free benchmark.

Theorem 7.1. *For downward-closed permutation environments, $\mathbf{E}[\text{IC}^{\text{RSEM}'}(\mathbf{v})] \geq \frac{1}{189} \text{EFO}^{(2)}(\mathbf{v})$ for all \mathbf{v} .*

Proof. Let \mathcal{B} denote the event that the market and sample are balanced. Lemma 7.2 states that the probability that the partition is balanced is at least:

$$\Pr[\mathcal{B}] \geq 0.339.$$

The expected IC revenue of RSEM' is at least its revenue conditioned on the partitioning into sample and market being balanced (denoted as event \mathcal{B}). I.e.,

$$\mathbf{E}[\text{IC}^{\text{RSEM}'}(\mathbf{v})] \geq \Pr[\mathcal{B}] \mathbf{E}[\text{IC}^{\text{RSEM}'}(\mathbf{v}) \mid \mathcal{B}].$$

Lemma 7.3 states that for any \mathbf{v} that balance implies the IC revenue of RSEM' is at least a $\frac{1}{32}$ fraction of the EF optimal revenue on the sample. Taking expectations,

$$\mathbf{E}[\text{IC}^{\text{RSEM}'}(\mathbf{v}) \mid \mathcal{B}] \geq \frac{1}{32} \mathbf{E}[\text{EFO}(\mathbf{v}_S) \mid \mathcal{B}].$$

Lemma 7.11 states that the EF optimal revenue on the sample is at least half the envy-free benchmark, in expectation and conditioned on a balanced partitioning. I.e.,

$$\mathbf{E}[\text{EFO}(\mathbf{v}_S) \mid \mathcal{B}] \geq \frac{1}{2} \text{EFO}^{(2)}(\mathbf{v}).$$

Combining the above inequalities we conclude that the IC revenue of RSEM' is at least a 189-approximation to the envy-free benchmark. I.e.,

$$\mathbf{E}[\text{IC}^{\text{RSEM}'}(\mathbf{v})] \geq \frac{1}{189} \text{EFO}^{(2)}(\mathbf{v}). \quad \square$$

7.1 Balanced Partitioning

We now show that with high probability the partitioning of the agents into the market and sample satisfies a natural balanced property. Recall that, by definition, agent 1 is in M . This balanced property is a double-sided version of the balanced property introduced by Feige et al. (2005).

Definition 7.2. *A partitioning (S, M) of agents $N = \{1, \dots, n\}$ is balanced if $1 \in M$ and $2 \in S$ and for any set of three or more of the highest valued agents both the market and sample contain at least a quarter of agents in the set. I.e., for $i \geq 3$, $|S \cap \{1, \dots, i\}| \geq i/4$ and $|M \cap \{1, \dots, i\}| \geq i/4$.*

Lemma 7.2. *Conditioning on $1 \in M$, a random partitioning (S, M) of N is balanced with probability at least 0.339.*

Proof. Conditioning on $1 \in M$ and $2 \in S$, the probability that either part is imbalanced can be calculated to be at most 0.161 by a simple probability of ruin analysis which comes from Feige et al. (2005) (details given below). By the union bound, both parts are balanced with probability at least 0.678. Agent 2 is in S with probability $1/2$ so the probability of balance conditioned on agent 1 in M is at least 0.339.

The following analysis from Feige et al. (2005) shows that the probability that S is imbalanced is at most 0.161. Consider the random variable $Z_i = 4|S \cap \{1, \dots, i\}| - i$; the balanced condition is equivalent to $Z_i \geq 0$ for all $i \geq 3$. By the conditioning $i = 2$ and $S \cap \{1, 2\} = \{2\}$ imply that $Z_2 = 2$. View Z_i as the positions of a random walk on the integers that starts from position two and takes three steps forward (at step i with $i \in S$) or one step back (at step i with $i \notin S$), each with probability one half. If this random walk ever arrives at position -1 the partitioning is imbalanced. This probability r of ever taking one step back in such a random walk can be calculated as the root of $r^4 - 2r + 1$ on interval $(0, 1)$ which is about 0.544. The probability of imbalance is then $r^3 \leq 0.161$ (i.e., if we ever take three steps back when starting from position two). By symmetry, the probability of imbalance in the market M is also at most 0.161. \square

7.2 Market revenue versus sample revenue

We now show that conditioned on a balanced partitioning of the agents into a market and sample, that the revenue of RSEM' from the market is close to the envy-free optimal revenue from the sample. The revenue of RSEM' is precisely $\text{IC}_M^S(\mathbf{v}_N)$, i.e., the revenue we get from the agents in M when using the virtual value functions from S and optimizing virtual values on the over the full set of agents N . We wish to compare this revenue to the envy-free optimal revenue on the sample, $\text{EFO}(\mathbf{v}_S)$. Our bound follows from three steps: the IC revenue of RSEM' from M is close to its EF revenue from M , the EF revenue of RSEM' from M is close to its EF revenue from the full set of agents N , and this EF revenue from the full set of agents is close to the optimal envy-free revenue from the sample.

Lemma 7.3. *Given a balanced partitioning (M, S) , $\text{IC}^{\text{RSEM}'}(\mathbf{v}) \geq \frac{1}{32} \text{EFO}(\mathbf{v}_S)$.*

Proof. The proof is given by the following sequence of inequalities:

$$\begin{aligned}
\text{IC}^{\text{RSEM}'}(\mathbf{v}) &= \text{IC}_M^S(\mathbf{v}_N) && \text{(Definition 7.1)} \\
&\geq \frac{1}{2} \text{EF}_M^S(\mathbf{v}_N) && \text{(Lemma 4.2)} \\
&\geq \frac{1}{8} \text{EF}_N^S(\mathbf{v}_N) && \text{(below, Lemma 7.4)} \\
&\geq \frac{1}{32} \text{EF}_S^S(\mathbf{v}_S) && \text{(below, Lemma 7.9)} \\
&= \frac{1}{32} \text{EFO}(\mathbf{v}_S). && \square
\end{aligned}$$

The first step (Lemma 4.2) of Lemma 7.3 was proven in Section 4. To see the second step notice that the envy-free payment of the agents $i \in N$ in $\text{EF}^S(\mathbf{v}_N)$ form a non-increasing sequence. This non-increasing sequence of payments can be plugged into the following lemma (Lemma 7.4) as the a_i s. The proof of the lemma follows directly from its statement.

Lemma 7.4. *Given a balanced partitioning (S, M) , for every non-increasing sequence a_1, \dots, a_n of nonnegative reals and all $i \in N$, $\sum_{j \in M \cap \{1, \dots, i\}} a_j \geq \frac{1}{4} \sum_{j \in \{1, \dots, i\}} a_j$.*

To show the third step, we need to give a detailed analysis of what happens in terms of envy-free revenue when we optimize for the wrong virtual values. To do that we will define and consider the *effective revenue curve*, $\tilde{\mathbf{R}}$, and *perceived revenue curve*, $\hat{\mathbf{R}}$. Intuitively, $\hat{\mathbf{R}}$ corresponds to the revenue we think we get when optimizing $\bar{\Phi}^S$ on \mathbf{v} , and $\tilde{\mathbf{R}}$ corresponds to the revenue curve we actually end up with.




Figure 4: Effective revenue curves and revenue dominance depicted.

Definition 7.3 (Effective revenue curve \tilde{R}). For values \mathbf{v} and ironed virtual valuations $\bar{\Phi}^S$ for S : group agents with equal nonnegative $\bar{\Phi}^S$ values into consecutive classes $\{1, \dots, n_1\}$, $\{n_1+1, \dots, n_2\}$, \dots , $\{n_{t-1}+1, \dots, n_t\}$ and define the effective revenue curve \tilde{R} from $R = R^{\mathbf{v}}$ by connecting the points $(0, 0)$, $(n_1, R(n_1))$, \dots , $(n_t, R(n_t))$ and then extending horizontally to $(n, R(n_t))$, i.e., ironing the values in each class.

Figure 4(a) depicts an example of the effective revenue curve. The three rays from the origin, which correspond to values at which $\bar{\Phi}^S$ makes a piece-wise jump, divide the first orthant into four regions. For every region, every point $(i, R(i))$ in the region (which corresponds to value v_i) has the same $\bar{\Phi}^S$ value. In each region these points get “ironed”, and hence the line segment in \tilde{R} .

Lemma 7.5. $EF^{\bar{\Phi}^S}(\mathbf{v}) = \sum_{i=1}^n \tilde{R}(i) \cdot (x_i^S(\mathbf{v}) - x_{i+1}^S(\mathbf{v}))$.

Proof.

$$\begin{aligned} EF^{\bar{\Phi}^S}(\mathbf{v}) &= \sum_{i=1}^n R(i) \cdot (x_i^S(\mathbf{v}) - x_{i+1}^S(\mathbf{v})) \\ &= \sum_{i=1}^n \tilde{R}(i) \cdot (x_i^S(\mathbf{v}) - x_{i+1}^S(\mathbf{v})) \end{aligned}$$

Here the first equality is by Lemma 3.5. To justify the second equality, note that whenever $\tilde{R}(i) \neq R(i)$, there are two cases: (1) i is in $\{n_{j-1}+1, \dots, n_j-1\}$ for some j , and so v_i and v_{i+1} have the same $\bar{\Phi}^S$ value, and hence $x_i^S(\mathbf{v}) = x_{i+1}^S(\mathbf{v})$; and (2) i is bigger than n_t , and so v_i and v_{i+1} both have negative $\bar{\Phi}^S$ value, and hence $x_i^S(\mathbf{v}) = x_{i+1}^S(\mathbf{v}) = 0$. \square

For a set of agents S , let \mathbf{v}_S denote $(\mathbf{v}_S, \mathbf{0}_{N-S})$, i.e., the valuation profile (of n agents) obtained from \mathbf{v} by decreasing the values of agents outside S to 0. Note that $\mathbf{v} = \mathbf{v}_N$. Let R^S and \tilde{R}^S be the revenue curve and ironed revenue curve of the valuation profile \mathbf{v}_S respectively.

Lemma 7.6. For all $1 \leq i \leq n$, $\tilde{R}(i) \geq \tilde{R}^S(i)$.

Proof sketch. Figure 4(b) depicts the relationship between the revenue curves. Observe that revenue curve R dominates R^S in the sense that for every slope t , the intersection of the ray $y = tx$ with R is farther away from the origin than its intersection with R^S . Transforming R and R^S to the effective revenue curves using the same ironed virtual valuation function $\bar{\Phi}^S$ do not change such dominance relationship, and moreover, because \tilde{R}^S is non-decreasing and concave, it follows that vertical dominance also holds, i.e., $\tilde{R}(i) \geq \tilde{R}^S(i)$ for all i . \square

Definition 7.4 (Perceived revenue curve \widehat{R}). *The perceived revenue curve for $\bar{\Phi}^S$ on \mathbf{v} is given by $\widehat{R}(i) = \sum_{j=1}^i \bar{\Phi}^S(v_j)$ for $i \in N$.*

Let $\hat{\mathbf{v}}$ be the valuation profile corresponding to \widehat{R} , i.e., $\hat{v}_i = \widehat{R}(i)/i$, and let $\mathbf{x}^{\hat{\mathbf{v}}}$ be the ironed virtual surplus maximizer for $\bar{\Phi}^{\hat{\mathbf{v}}}$.

Lemma 7.7. $x_i^S(\mathbf{v}) = x_i^{\hat{\mathbf{v}}}(\hat{\mathbf{v}})$.

Proof. Compare running the ironed virtual surplus maximizer \mathbf{x}^S for $\bar{\Phi}^S$ on \mathbf{v} with running $\mathbf{x}^{\hat{\mathbf{v}}}$ for $\bar{\Phi}^{\hat{\mathbf{v}}}$ on $\hat{\mathbf{v}}$, the ironed virtual valuation of agent i in either case is equal to $\bar{\Phi}^S(v_i)$. Therefore these two ironed virtual surplus optimizers will choose the same allocation, and the lemma follows. \square

Lemma 7.8. *Given a balanced partitioning (S, M) , then $\bar{R}^S(i) \geq \frac{1}{4}\widehat{R}(i) \geq \frac{1}{4}\bar{R}^S(i)$ for all $1 \leq i \leq n$.*

Proof. For each i , $\widehat{R}(i)$ is the sum of i largest ironed virtual values in N with respect to $\bar{\Phi}^S$ while $\bar{R}^S(i)$ is the sum of the i largest with respect to S . Therefore $\widehat{R} \geq \bar{R}^S(i)$. Since (S, M) is double-side balanced, applying Lemma 7.4, we also have that for all i , $\bar{R}^S(i) \geq \frac{1}{4}\widehat{R}(i)$. \square

Now we are ready prove the following key lemma:

Lemma 7.9. *For any downward-closed permutation environments, any valuation profile \mathbf{v} , and balanced partitioning (S, M) , $\text{EF}^{\bar{\Phi}^S}(\mathbf{v}_N) \geq \frac{1}{4}\text{EF}^{\bar{\Phi}^S}(\mathbf{v}_S) = \frac{1}{4}\text{EFO}(\mathbf{v}_S)$.*

Proof. Let \mathbf{x}^S and $\mathbf{x}^{\hat{\mathbf{v}}}$ be short-hands for the ironed virtual surplus optimizers with ironed virtual valuation functions defined for \mathbf{v}_S and $\hat{\mathbf{v}}$, respectively. The proof is by the following inequalities:

$$\begin{aligned} \text{EF}^{\bar{\Phi}^S}(\mathbf{v}_N) &= \sum_i \widetilde{R}(i) \cdot (x_i^S(\mathbf{v}_N) - x_{i+1}^S(\mathbf{v}_N)) \\ &= \sum_i \widetilde{R}(i) \cdot (x_i^{\hat{\mathbf{v}}}(\hat{\mathbf{v}}) - x_{i+1}^{\hat{\mathbf{v}}}(\hat{\mathbf{v}})) \\ &\geq \frac{1}{4} \cdot \sum_i \widehat{R}(i) \cdot (x_i^{\hat{\mathbf{v}}}(\hat{\mathbf{v}}) - x_{i+1}^{\hat{\mathbf{v}}}(\hat{\mathbf{v}})) \\ &\geq \frac{1}{4} \cdot \sum_i \widehat{R}(i) \cdot (x_i^S(\mathbf{v}_S) - x_{i+1}^S(\mathbf{v}_S)) \\ &\geq \frac{1}{4} \cdot \sum_i \bar{R}^S(i) \cdot (x_i^S(\mathbf{v}_S) - x_{i+1}^S(\mathbf{v}_S)). \end{aligned}$$

Here the first two equalities are guaranteed by our definitions of \widetilde{R} and \widehat{R} . The first inequality is by Lemma 7.6 and Lemma 7.8, the second inequality is by the optimality of $\mathbf{x}^{\hat{\mathbf{v}}}$ for $\hat{\mathbf{v}}$, and the third inequality is by Lemma 7.8 again. \square

7.3 Expected sample revenue versus the envy-free benchmark

We now show that the expected envy-free revenue of the sample compares favorably with the envy-free benchmark; this is the last ingredient in the proof of Theorem 7.1. We will make this argument conditioned on a balanced partitioning; however, the result is true for any symmetric conditioning (including none at all).

Lemma 7.10. *For a partitioning (S, M) of N , we have that $\text{EFO}(\mathbf{v}_S) + \text{EFO}(\mathbf{v}_M) \geq \text{EFO}(\mathbf{v}_N)$.*

Proof. $\text{EFO}(\mathbf{v}_N) = \text{EFO}(\mathbf{v}_{S \cup M})$ is the maximum revenue we can get from $S \cup M$ subject to the envy free constraints. Let agents in M contribute total revenue R to $\text{EFO}(\mathbf{v}_N)$. By setting the agents in S to have zero valuations to obtain valuation profile \mathbf{v}_S , we basically removed envy-freedom constraints between agents in S and agents in M . With fewer envy-freedom constraints, the maximum envy-free revenue we can get from M , i.e., $\text{EFO}(\mathbf{v}_M)$, can only be larger. Similarly, the total revenue that S contributes to $\text{EFO}(\mathbf{v}_N)$ is at most $\text{EFO}(\mathbf{v}_S)$, and our lemma follows. \square

Lemma 7.11. $\mathbb{E}[\text{EFO}(\mathbf{v}_S) \mid \mathcal{B}] \geq \frac{1}{2} \text{EFO}^{(2)}(\mathbf{v})$ where \mathcal{B} is the event that S and M are balanced.

Proof. The lemma follows from Lemma 7.10, symmetry, and the fact that agent 1 is always in M . \square

References

- Aggarwal, G. and Hartline, J. (2006). Knapsack auctions. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms*.
- Alaei, S., Malekian, A., and Srinivasan, A. (2009). On random sampling auctions for digital goods. In *ACM Conference on Electronic Commerce*, pages 187–196.
- Balcan, M.-F., Blum, A., Hartline, J. D., and Mansour, Y. (2008). Reducing mechanism design to algorithm design via machine learning. *J. Comput. Syst. Sci.*, 74(8):1245–1270.
- Baliga, S. and Vohra, R. (2003). Market research and market design. *Advances in Theoretical Economics*, 3.
- Bikhchandani, S., de Vries, S., Schummer, J., and Vohra, R. (2008). Ascending auctions for integral (poly)matroids with concave nondecreasing separable values. In *Proc. 19th ACM Symp. on Discrete Algorithms (SODA)*, pages 864–873.
- Bulow, J. and Klemperer, P. (1996). Auctions versus negotiations. *American Economic Review*, 86(1):180–194.
- Bulow, J. and Roberts, J. (1989). The simple economics of optimal auctions. *The Journal of Political Economy*, 97:1060–90.
- Clarke, E. H. (1971). Multipart pricing of public goods. *Public Choice*, 11:17–33.
- Dughmi, S., Roughgarden, T., and Sundararajan, M. (2009). Revenue submodularity. In *EC ’09: Proceedings of the tenth ACM conference on Electronic commerce*, pages 243–252.
- Edelman, B., Ostrovsky, M., and Schwarz, M. (2007). Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *The American Economic Review*, 97(1):242–259.
- Feige, U., Flaxman, A., Hartline, J., and Kleinberg, R. (2005). On the Competitive Ratio of the Random Sampling Auction. In *Proc. 1st Workshop on Internet and Network Economics*, pages 878–886.

- Goldberg, A., Hartline, J., Karlin, A., Saks, M., and Wright, A. (2006a). Competitive auctions. *Games and Economic Behavior*.
- Goldberg, A., Hartline, J., Karlin, A., Saks, M., and Wright, A. (2006b). Competitive auctions. *Games and Economic Behavior*, 55(2):242–269.
- Goldberg, A. V. and Hartline, J. D. (2003). Competitiveness via consensus. In *Proc. 14th ACM Symp. on Discrete Algorithms (SODA)*. ACM/SIAM.
- Goldberg, A. V., Hartline, J. D., and Wright, A. (2001). Competitive auctions and digital goods. In *Proc. 12th ACM Symp. on Discrete Algorithms (SODA)*, pages 735–744.
- Groves, T. (1973). Incentives in teams. *Econometrica*, 41:617–631.
- Ha, B. and Hartline, J. (2011). Biased random sampling and profit extraction. Manuscript.
- Ha, B. and Hartline, J. (2012). Mechanism design via profit extraction, consensus, and cross-check. In *Proc. 23rd ACM Symp. on Discrete Algorithms (SODA)*.
- Hartline, J. (2011). Approximation in mechanism design. Manuscript.
- Hartline, J. and McGrew, R. (2005). From optimal limited to unlimited supply auctions. In *Proc. 7th ACM Conf. on Electronic Commerce (EC)*, pages 175–182.
- Hartline, J. and Roughgarden, T. (2008). Optimal mechanism design and money burning. In *Proc. 39th ACM Symp. on Theory of Computing (STOC)*, pages 75–84.
- Hartline, J. and Roughgarden, T. (2009). Simple versus optimal mechanisms. In *Proc. 11th ACM Conf. on Electronic Commerce (EC)*.
- Ichiba, T. and Iwama, K. (2010). Averaging techniques for competitive auctions. In *Proceedings of the Seventh Workshop on Analytic Algorithmics and Combinatorics*, pages 74–81.
- Jackson, M. and Kremer, I. (2007). Envy-freeness and implementation in large economies. *Review of Economic Design*, 11:185–198.
- Lehmann, D. J., O’Callaghan, L., and Shoham, Y. (2002). Truth revelation in approximately efficient combinatorial auctions. *J. ACM*, 49(5):577–602.
- Myerson, R. (1981). Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73.
- Oxley, J. G. (1992). *Matroid Theory*. Oxford Science Publications. The Clarendon Press, Oxford University Press, New York.
- Rado, R. (1952). An equality. *J. London Math. Soc.*, 27:1–6.
- Segal, I. (2003). Optimal pricing mechanisms with unknown demand. In *American Economic Review* 93.
- Talwar, K. (2003). The price of truth: Frugality in truthful mechanisms. In *Proc. 20th Annual Symp. on Theoretical Aspects of Computer Science (STACS)*, pages 608–619.

- Varian, H. (2007). Position auctions. *International Journal of Industrial Organization*, 25(6):1163–1178.
- Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *J. of Finance*, 16:8–37.
- Wilson, R. (1987). Game theoretic analysis of trading processes. *Advances in Economic Theory*.

A Proofs

A.1 Proof of Lemma 5.4

Fix some number m .

The matroid we use is a partition matroid. In general in a partition matroid, the ground set is partitioned into a number of disjoint sets, or sectors, where each sector is associated with a capacity number. A set is feasible if and only if its intersection with each sector does not exceed the capacity number of the sector.

Now we define the partition matroid we use. For each $k \in \{1, \dots, m\}$, a type k sector contains m^{3k-1} elements or agents and has capacity one. There are m^{2m-2k} disjoint type k sectors for each $k \in \{1, \dots, m\}$. So total number of agents n is at most $m^{O(m)}$. Hence m is at least of order $\frac{\log n}{\log \log n}$.

Next we define the “sydney opera house distribution”. The distribution F is such that the value is distributed according to uniform distribution $[m^{2k+1} - \epsilon, m^{2k+1} + \epsilon]$ with probability $\frac{1}{m^{3k}} - \frac{1}{m^{3k+3}}$ for $k \in \{0, \dots, m-1\}$, and with probability $\frac{1}{m^{3k}}$ for $k = m$. Here we take ϵ to be some sufficiently negligible positive amount, and we will often omit ϵ related terms. So for each k the revenue function R at $\frac{1}{m^{3k+3}}$ has left limit $R(\frac{1}{m^{3k+3}}-) = \frac{m^{2k+3}}{m^{3k+3}} = \frac{1}{m^k}$, and right limit $R(\frac{1}{m^{3k+3}}+) = \frac{m^{2k+1}}{m^{3k+3}} = \frac{1}{m^{k+2}}$. Hence the ironed virtual valuation between quantile $\frac{1}{m^{3k+3}}$ to quantile $\frac{1}{m^{3k}}$ is $\frac{\frac{1}{m^{k-1}} - \frac{1}{m^k}}{\frac{1}{m^{3k}} - \frac{1}{m^{3k+3}}} \approx m^{2k+1}$. Note that the ironed virtual valuation is equal to valuation, ignoring minor terms.

To calculate the revenue of Myerson’s auction, for a type k sector, there are m^{3k-1} agents. With probability at least $1 - (1 - \frac{1}{m^{3k}})^{m^{3k-1}} \approx \frac{1}{m}$, the highest agent is in quantile range $(0, \frac{1}{m^{3k}})$, with ironed virtual valuation at least m^{2k+1} . So the expected ironed virtual valuation from a type k sector is at least m^{2k} . Multiplied by the number of type k sectors the total ironed virtual valuation, and hence expected revenue is at least $\sum_k m^{2k} \cdot m^{2m-2k} = m \cdot m^{2m}$.

To calculate the revenue of VCG with some reserve r , suppose w.l.o.g. $r \approx m^{2k'+1}$ for some k' . For a type k sector with m^{3k-1} agents, the dominant amount of revenue is obtained from the following two cases:

1. When there are at least two agents with value at least $m^{2k+1} - \epsilon$ (i.e. in quantile $\frac{1}{m^{3k}}$), the lower of which has value at most $m^{2k+1} + \epsilon$. This happens with probability roughly $\frac{1}{m^2}$, and gives revenue m^{2k+1} . Therefore the expected revenue we get from this case is $\frac{1}{m^2} \cdot m^{2k+1}$, which multiplied by the number of type k sector, is $O(m^{2m-1})$.
2. When $k = k'$, and there is at least one agent who beats the reserve $m^{2k'+1}$. This happens with probability at most $\frac{1}{m}$. Therefore the expected revenue from this case is $m^{2k'+1} \cdot \frac{1}{m}$, which multiplied by the number of type k sectors is $O(m^{2m})$.

Summing over all k , the total expected revenue of VCG with reserve r is at most $O(m^{2m}) + m \cdot O(m^{2m-1}) = O(m^{2m})$, which is less than that of Myerson's auction by a factor of $\Omega(m) = \Omega(\log n / \log \log n)$.

A.2 Proof of Lemma 4.4

Let there be $n + 1$ agents. The “1 vs n ” set system has two maximum feasible sets, one is a singleton set and the other one has size n . These two sets are disjoint. We define the valuation profile by specifying the virtual valuations. There are n “small” agents with virtual values $v + \epsilon, v + 2\epsilon, \dots, v + n\epsilon$ respectively, and one “big” agent with virtual value $nv + \frac{n(n+1)}{2}\epsilon - \epsilon^2$ for some small positive ϵ . The choice of the ϵ terms is such that for the sum of the virtual valuations of the first n agents to beat the big agent, no small agent can lower her virtual value to some other agent's virtual value. We will ignore ϵ terms from now on. Correspondingly, one can calculate the revenue curve, and then derive the valuations of the agents: the valuation of the big agent is nv , and the small agents have values $\frac{n+1}{2}v, \frac{n+2}{3}v, \dots, \frac{2n}{n+1}v$, ignoring ϵ terms. The allocation rule is the ironed virtual surplus optimizer w.r.t. this valuation profile. Note that a reserve of $\frac{2n}{n+1}v$ is set because any value lower than this corresponds to a negative ironed virtual value.

Observe that every agent wins if and only if she is assigned to the size n set, which happens with probability $n/(n+1)$. Therefore the EF revenue is $\frac{2n}{n+1}v \cdot \frac{n}{n+1} \cdot (n+1) = \frac{2n^2}{n+1}v$. To calculate the IC revenue, with probability $n/(n+1)$, the big agent is assigned to the size n set, and every of the n winning agents pays the reserve $\frac{2n}{n+1}v$. Also with probability $1/(n+1)$, the big agent is assigned to the singleton set, and every agent has to pay her own value, which sums up to $\Theta(nv \log(n))$. Therefore the IC revenue is $\frac{2n}{n+1}v \cdot \frac{n}{n+1} \cdot n + \frac{1}{n+1}\Theta(nv \log(n))$, which is larger than EF revenue for sufficiently large n .

A.3 Proof of Lemma 5.2

Proof. By the reduction from matroid permutation environments to multi-unit environments, it is sufficient to prove the statement for k -unit auctions. Let $\bar{\Phi}$ be the ironed virtual surplus maximizer for F . We first upper-bound IC revenue from bidders 2 to n :

$$\begin{aligned} \text{IC}_{2\dots n}^F(\mathbf{v}) &\leq \text{EF}_{2\dots n}^{\bar{\Phi}}(\mathbf{v}) \\ &\leq \text{EFO}(v_2, v_3, \dots, v_n, 0) \\ &\leq \text{EFO}(\mathbf{v}^{(2)}) \\ &= \text{EFO}^{(2)}(\mathbf{v}) \end{aligned}$$

Here the first inequality is by Lemma 4.3, and the last equality is by definition of $\text{EFO}^{(2)}$.

To see the second inequality, both left hand side and right hand side correspond to the maximum envy-free revenue from $2 \dots n$ that correspond to some outcome with at most k items allocated and no envy among $2 \dots n$, except that in the right hand side, the outcome that maximizes this revenue is chosen.

To see the third inequality, note that the revenue curve of $\mathbf{v}^{(2)}$ dominates that of $(v_2, v_3, \dots, v_n, 0)$, and hence by Lemma 3.2, for every allocation, the envy-free revenue for $\mathbf{v}^{(2)}$ can only be higher.

Next we upper-bound IC revenue from bidder 1, where the expectation is over i.i.d. draws from

distribution F .

$$\begin{aligned} E[\text{IC}_1^{\bar{\Phi}}(\mathbf{v})] &\leq E[\text{IC}^{\bar{\Phi}}(\mathbf{v}) \text{ for single item auction}] \\ &\leq 2E[v_2] \\ &\leq 2E[\text{EFO}^{(2)}(\mathbf{v})] \end{aligned}$$

Here the second inequality is by the tail regularity assumption. The third inequality is because $\text{EFO}^{(2)}(\mathbf{v}) \geq v_2$.

To see the first inequality, consider the mechanism that first runs $\bar{\Phi}$, and then only allows the highest bidder (bidder 1) to win. The IC payment of bidder 1 is the maximum of the second highest bid and the threshold bid for bidder 1 to win in $\bar{\Phi}$. This is as much as the threshold (or revenue) from bidder 1 in $\bar{\Phi}$ as in the left hand side. On the other hand, the IC revenue of this mechanism is at most that of the optimal single-item auction, which equals to the right hand side.

Together, we have that $E[\text{IC}^{\bar{\Phi}}(\mathbf{v})] \leq 3E[\text{EFO}^{(2)}(\mathbf{v})]$.

□